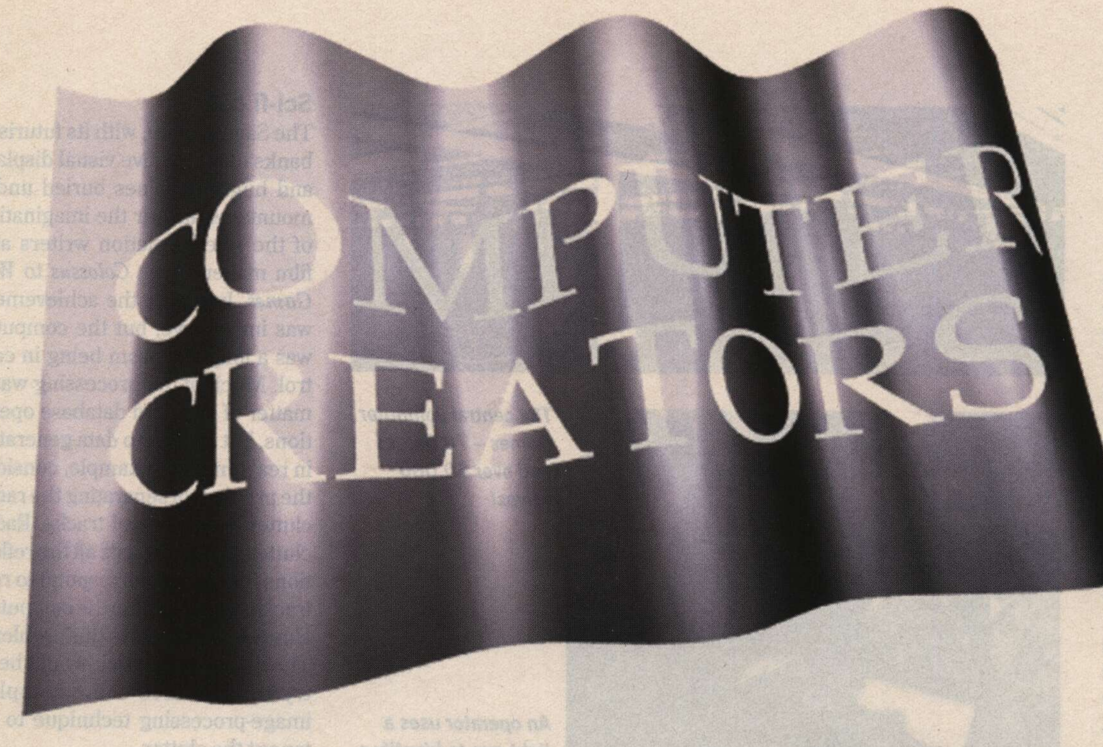

COMPUTER CREATORS

Mike James looks at Sage – technology which developed out of the threat of nuclear war

It is a common belief that war drives the pace of technology. In the case of computing this is certainly true. The majority of early machines, the ENIAC and the Mark I for example, were built for military calculations. However, in the case of computing it is arguable that it was the Cold War that produced the greatest advances. The computer is an ideal instrument for gathering and processing the sort of information that you need to keep an eye on what your potential enemies are doing. The Sage (semi-automatic ground environment) system was the first real-time computer and it was huge – deep inside caverns, enormous computers tracked every movement in the sky and protected the USA from sneak mass nuclear attack. The mythology makes Sage sound like some early version of the Star Wars program, but in fact the reality is much more interesting.

File card: Sage

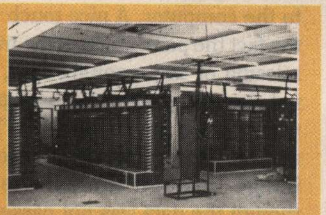

Dates: 1950 – 1983

Achievements:

- **First use of core memory in a production machine**
- **First duplex machine system for high reliability**
- **Development of modems for phone line data transmission**
- **Early time-sharing system**
- **Use of DMA for peripheral data transfers**
- **Invention of light gun and interactive display**

THE WHIRLWIND COMPUTER was built at MIT and in the 1940s it was the fastest computer available, but it was going nowhere. The Whirlwind was a real-time computer suitable for flight simulators and the like, but at the time few people understood the potential of such a machine. Then the Russians exploded a nuclear device and the USA was shocked out of its complacency. The Cold War had started.

The US had to face the prospect of nuclear missiles entering its air space and remaining undetected until they gave themselves away in the last split microsecond of their existence. The Air Defense Engineering Committee was formed in 1949 and decided that what was needed was a tracking system that combined the talents of the newly invented computer with those of radar. The radar system could be used to detect airborne objects and the computer could keep track of them all and make sure that every moving object was a known object.

This was a brave concept given that both computers and radar were new technologies and nothing like what they proposed had even remotely been shown to work. A network of radar stations communicating with a central computer data processing station really looks ambitious when you realise that at the time no one had even invented the modem, a reliable method of mass storage, a graphical visual display unit, and computers were still built using valves!

A research lab was founded in 1951 to create the new system. At first this was known as Project Lincoln and then the Lincoln Laboratory, which drew on developments from other research centres. In particular MIT supplied the computer hardware in the form of Whirlwind I (see last month's *Computer Creators*).

The data communications came from the Air Force Cambridge Research Laboratory (CRL). At this early stage researchers estimated that they could do the job with just 8K of memory and only a few thousand lines of code. The true magnitude of the problem only sank in slowly. If they had started out with a complete systems analysis the task would probably have been declared impossible.

As it happened the researchers chose to adopt a step-by-step approach, solving each problem as it came up. Sometimes their solutions were lower-tech than we might expect today, but in fact they often involved pushing the limits of the current technology well beyond anything that had been achieved before.

The time-share option

Before committing to a full system, an experimental implementation called the Cape Cod system (after its location) was tried out. The Cape Cod system was the first to use computer time-sharing, a CRT display system and light pens. This was not only a forerunner of the defence system to come, but the real-time GUI interfaces that we all use today. Cape Cod was a test bench for the Whirlwind computer, the software, the radar installations and the communication techniques needed to link them all together. Telephone lines were leased and data from remote radar stations was transmitted to the Whirlwind using the first modems.

At first, results were poor. Crosstalk and dialling noise generated multiple false targets. The data transmission worked only 50 percent of the time, but at least it worked. To test the entire system researchers enlisted the help of the local Air National Guard who flew a pair of small propeller-driven planes up and down the coast. In early 1951 the Guard succeeded in demonstrating live the ability to issue intercept instructions with an

➧

error of only 1,000 yards. Impressive even if the planes were only moving at one quarter of the speed of a typical jet.

Enter IBM

The Cape Cod system used the Whirlwind computer, but for the production system something more manufacturable was needed. This is where IBM comes in. IBM and MIT co-operated on a commercial development of the Whirlwind – essentially the Whirlwind II, but the production version was known as the FSQ-7. It is clear that IBM's involvement had a spin-off into its early range of machines pre-dating the System 360 project.

The prototype machine – the XD-1 – had 8,000 words of core memory and this immediately proved too small because of the performance degradation produced by having to page to magnetic drum. Jay Forrester met the requirement for more memory by producing a 65,000-word core memory – a staggeringly large amount of high-speed memory for the time. The machine was as reliable as any valve machine could be, but it still wasn't good enough.

The solution was to use two machines in duplex arrangement. When one machine was processing radar data the other would be used for routine housekeeping and preventative maintenance. There were two problems with this idea. The first was that the Air Force had swallowed the bill for one set of computers, but now this was to be doubled! The second was that for the duplex system to work, a way of allowing one machine to hot swap with the other was needed.

That is, just having two machines available is no use when continuity of data processing is all important. At first researchers thought in terms of getting both computers to do everything all of the time, but this would be inefficient. In the end the backup computer did some of the tasks while keeping up to date via a shared magnetic drum. The duplication of the main machines was only part of the problem – all critical peripherals also had to be copied and a way to select which machine controlled which peripheral organised .

The radar centres had to transmit their data back to the command centres for processing. To achieve this the modem had to be invented. At first no one could understand why it wasn't working. They asked Bell for a specification its telephone lines conformed to and were surprised by the answer – "The lines are all right if you can talk over them." Eventually the problem was solved by asking for a 300-mile telephone loop, injecting signals into one end and looking at what came out of the other. As a result, all the modems were redesigned and made to work at 1,300 baud – as long as the telephone lines were specially conditioned.

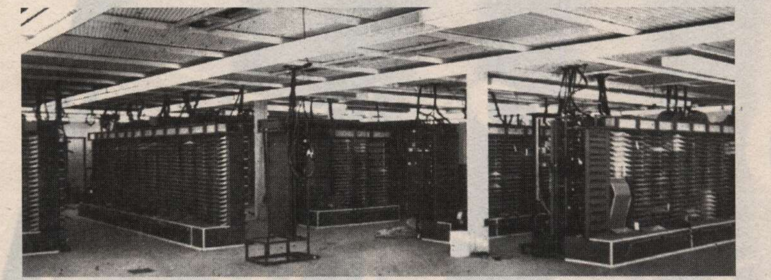

The central computer frames – 70 frames and over 60,000 valves!

An operator uses a light gun to identify a track on the Charactron tube

Going underground

Of course, being a military installation and an important potential target for nuclear attack, the command centres needed to be protected. But how to protect a machine the size of a building and one that was still evolving? The first idea as to bury the whole thing underground. Digging started on the first Direction Centre only to hit the water table, test borings for the second site hit solid granite. From then on the Centre grew upwards floor by floor. Others were buried in mountainsides and some were built like huge concrete block houses. All had their own dual generator sets, cooling plants and a single floor dedicated to the FSQ-70 computer – consisting of 70 frames of 60,000 valves.

Each machine weighed 250 tons and needed a 3,000 kilowatt power supply. The final version of the machine was a single address binary computer with a 32-bit word size. The core memory was 8,192 words with a cycle time of six micro seconds (ie 0.17MHz). In practice it managed a very respectable 75,000 32-bit operations per second.

Another technique pioneered was the use of cycle-stealing I/O channels – today we would probably call this DMA. Although it was a general purpose computer it had some special features to enable it to work with x,y co-ordinates. Its registers could treat a single 32-bit word as a pair of 16-bit values – effectively doubling its throughput. A second level of storage was provided by 12 magnetic drums holding 12,000 32-bit words each. These were also used to buffer incoming data until it could be processed.

The results were displayed on specially designed video consoles. These used large, 19" Charactron tubes which allowed graphics and text to be mixed. The CRT tube contained a mask which held the shape of each character. The beam was deflected to pass through one of the character shapes and then steered to a specific screen location. There was also a small, 5" text-only display used to give details such as weather and other background information.

The operator could interact with the screen with a remarkable new device – a light gun. This gun was pointed at the screen and a trigger pulled to select a radar track. Once selected the operator could type in information that identified the type of track.

Clearly the operator played an important role in processing the data and there is no suggestion that the Sage system represented a fully automatic defence system. However, to the casual user it looked as if the day of automatic war had arrived.

Sci-fi hell

The Sage system, with its futuristic banks of interactive visual displays and huge machines buried under mountains did stir the imagination of the science fiction writers and film makers from *Colossus* to *War Games*. In reality the achievement was impressive, but the computer was a long way from being in control. Much of the processing was a matter of standard database operations, but applied to data generated in real time. For example, consider the problem of separating the radar clutter from the real tracks. Radar clutter, in other words all the reflections that do not correspond to real tracks, would swamp the computer. How would you solve this problem? The modern approach would be to try to introduce some complex image-processing technique to filter out the clutter.

The Sage team didn't resort to the full high-tech, but it did find a solution. The radar image was displayed on a PPI (plan position indicator – ie a standard radar display) and a special paint coating was applied to cover up the stationary clutter. An operator actually painted out the coastline and any other regular reflections on the tube. A photocell was fitted to pick up the flashes that were not masked. The pulses from the photocell were used to gate the data being passed to the computer. A nice touch was the fact the paint stopped the transmission of the scanning spot, but not the afterglow of the phosphor which was a different colour. This allowed the operators to still see the image through the paint even though the data was not passed to the computer. That way there was no possibility of the Russians sneaking up through the clutter!

The full Sage system used 24 FSQ-7 machines. The first Direction Centre was deployed in 1958 and the final Direction Centre was opened in 1963. The last one was closed down in 1983 – representing some of the oldest computers in use at the time.

The management structure that created it seems to have been very loose to say the least. Budgets were overrun without too much worry, engineers took implementation decisions without reference to a committee and a sense of urgency and national security drove the design teams. Further away from the Second World War and with the management techniques of today it is doubtful if 10 percent of the results could have been achieved with 200 percent of the resources! ■